

REPORT DOCUMENTATION PAGE			Form Approved OMB NO. 0704-0188		
<p>The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA, 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.</p> <p>PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.</p>					
1. REPORT DATE (DD-MM-YYYY) 25-08-2011		2. REPORT TYPE Conference Proceeding		3. DATES COVERED (From - To) -	
4. TITLE AND SUBTITLE Detection of explosive hazards using spectrum features from forward-looking ground penetrating radar imagery			5a. CONTRACT NUMBER W911NF-10-1-0279		
			5b. GRANT NUMBER		
			5c. PROGRAM ELEMENT NUMBER 633606		
6. AUTHORS Justin Farrell, Timothy C. Havens, K. C. Ho, James M. Keller, Tuan T. Ton, David C. Wong, Mehrdad Soumekh			5d. PROJECT NUMBER		
			5e. TASK NUMBER		
			5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAMES AND ADDRESSES University of Missouri - Columbia Office of Sponsored Programs The Curators of the University of Missouri Columbia, MO 65211 -			8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) U.S. Army Research Office P.O. Box 12211 Research Triangle Park, NC 27709-2211			10. SPONSOR/MONITOR'S ACRONYM(S) ARO		
			11. SPONSOR/MONITOR'S REPORT NUMBER(S) 57940-EL.2		
12. DISTRIBUTION AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.					
13. SUPPLEMENTARY NOTES The views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy or decision, unless so designated by other documentation.					
14. ABSTRACT Buried explosives have proven to be a challenging problem for which ground penetrating radar (GPR) has shown to be effective. This paper discusses an explosive hazard detection algorithm for forward looking GPR (FLGPR). The proposed algorithm uses the fast Fourier transform (FFT) to obtain spectral features of anomalies in the FLGPR imagery. Results show that the spectral characteristics of explosive hazards differ from that of background clutter and are useful for rejecting false alarms (FAs). A genetic algorithm (GA) is developed in order to select a subset of					
15. SUBJECT TERMS explosive hazards,spectrum features, forward-looking ground penetrating radar					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	15. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON James Keller
a. REPORT UU	b. ABSTRACT UU	c. THIS PAGE UU			19b. TELEPHONE NUMBER 573-882-7339

Report Title

Detection of explosive hazards using spectrum features from forward-looking ground penetrating radar imagery

ABSTRACT

Buried explosives have proven to be a challenging problem for which ground penetrating radar (GPR) has shown to be effective. This paper discusses an explosive hazard detection algorithm for forward looking GPR (FLGPR). The proposed algorithm uses the fast Fourier transform (FFT) to obtain spectral features of anomalies in the FLGPR imagery. Results show that the spectral characteristics of explosive hazards differ from that of background clutter and are useful for rejecting false alarms (FAs). A genetic algorithm (GA) is developed in order to select a subset of spectral features to produce a more generalized classifier. Furthermore, a GA-based K-Nearest Neighbor probability density estimator is employed in which targets and false alarms are used as training data to produce a two-class classifier. The experimental results of this paper use data collected by the US Army and show the effectiveness of spectrum based features in the detection of explosive hazards.

Conference Name: Detection and Sensing of Mines, Explosive Objects, and Obscured Targets XVI

Conference Date:

Detection of Explosive Hazards Using Spectrum Features From Forward-Looking Ground Penetrating Radar Imagery

Justin Farrell^a, Timothy C. Havens^a, K.C. Ho^a, James M. Keller^a,
Tuan T. Ton^b, David C. Wong^b, and Mehrdad Soumekh^c

^aElectrical and Computer Engineering, University of Missouri, Columbia, MO, 65211;

^bU.S. Army Night Vision & Electronic Sensors Directorate, Fort Belvoir, Virginia, 22060;

^cElectrical Engineering, University of New York at Buffalo, Amherst, NY, 14260

ABSTRACT

Buried explosives have proven to be a challenging problem for which *ground penetrating radar* (GPR) has shown to be effective. This paper discusses an explosive hazard detection algorithm for *forward looking* GPR (FLGPR). The proposed algorithm uses the *fast Fourier transform* (FFT) to obtain spectral features of anomalies in the FLGPR imagery. Results show that the spectral characteristics of explosive hazards differ from that of background clutter and are useful for rejecting *false alarms* (FAs). A *genetic algorithm* (GA) is developed in order to select a subset of spectral features to produce a more generalized classifier. Furthermore, a GA-based K-Nearest Neighbor probability density estimator is employed in which targets and false alarms are used as training data to produce a two-class classifier. The experimental results of this paper use data collected by the US Army and show the effectiveness of spectrum based features in the detection of explosive hazards.

1. INTRODUCTION

The threat posed by buried explosives is a challenging problem to which many solutions have been proposed. Research on downward-looking *ground penetrating radar* (GPR) has shown impressive results in recent years [1]. *Forward-looking* GPR (FLGPR) has not yet demonstrated the same performance, but the potential to detect explosives hazards before they are encountered is of special interest in applications such as a vehicle traveling a road in which mines are buried. The data used for this paper was collected in this way. A system called ALARIC, developed by the U.S. Army Night Vision & Electronic Sensors Directorate (NVESD), travels a road with buried targets of known type and location as the FLGPR collects data. Targets vary in terms of metal content (high to low) and buried depth (1 to 6 inches). Unfortunately, FLGPR is susceptible to objects other than mines which must be addressed in the detection process. This can produce *false alarms* (FAs), which need to be mitigated without hindering the detection of mines. FA rejection can be done using algorithms that discriminate between background clutter and points of interest based on features that characterize both FAs and mines.

Figure 1 illustrates the steps of our proposed algorithm to detect and classify anomalies in FLGPR. This paper discusses the extraction of spectral features, a classifier and a *genetic algorithm* (GA) for feature selection. In section 2, an overview of previous work is presented. A pre-screener is discussed in section 3 and section 4 is a discussion of different 1D and 2D spectral features. Section 5 is the presentation of the classifier and the GA for feature selection. Furthermore, in section 6 we discuss results.

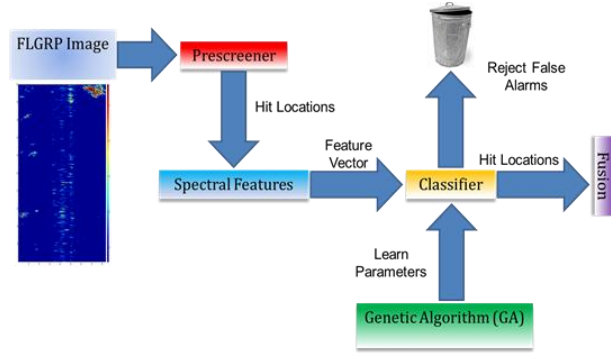


Figure 1. Proposed processing framework.

2. RELATED WORK

In [2], a one-class classifier was developed to detect explosive hazards in FLGPR. The classifier was trained by calculating a multivariate normal distribution from feature values of FAs from training data. These features were obtained from the spectral analysis of the FLGPR image, while the distribution parameters were calculated using the maximum likelihood estimator. The Mahalanobis-metric was used to determine how well a new test point matched the distribution described by the FAs. Since the classifier models FAs, a large Mahalanobis distance indicates dissimilarity between the test point and the model, providing evidence that the test point is a positive detection. Conversely, if the Mahalanobis distance is small, a test point is classified as a FA. This requires that a threshold be chosen such that anything greater than the threshold be classified as a positive detection and anything less than the threshold be classified as a FA. This threshold along with the classifier parameters and a reduced feature set was found using an exhaustive search. The search was conducted with the purpose of maintaining 100% probability of detection (PD) while reducing FAs. This method of classification did improve the FA rate significantly. However, the training of the classifier was done with only the FAs. This paper extends this prior work and proposes a two-class classifier which trains on both the FAs and the positive detections.

3. PRE-SCREENER

The pre-screener was originally proposed in [3]. The intention of the pre-screener is to provide a list of possible hit locations that are then sent to a classifier. To do this, a method to find areas of interest is necessary. A simple threshold could be used to find hits, however this generally results in a significant number of FAs and, in many cases, actual targets go undetected. Our approach uses a *locally-adaptive standard deviation filter* (LASDF) as well as a maximum order filter to find hit locations. The LASDF calculates the standard deviation in a rectangular halo around each pixel. Figure 2 shows a radar sub-image along with the halo used in the computation of the standard deviation.

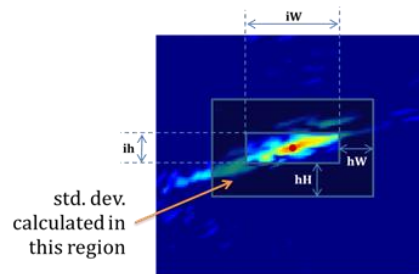


Figure 2. Illustration of the locally-adaptive standard deviation filter.

Each pixel in the FLGPR image is divided by the local standard deviation of the region depicted in Figure 2. This procedure results in a filtered image, $G_f(u, v)$, which is then processed further with a local-maxima finding algorithm. This algorithm first uses a maximum order filter with a 3m x 1.5m kernel. Figure 3 shows an example of the filtered image alongside the LASDF image.




Figure 3. Maximum order-filtered images of FLGPR images – target locations indicated by white circles.

Hits are found in the order filtered image according to

$$A = \arg_{\{u,v\}} \{G_f(u, v) \geq \min\{O_f(u, v), -60\}\}.$$

Where A is the set of local-maxima locations, $G_f(u, v)$ is the LASDF image and $O_f(u, v)$ is the maximum order-filter image. The lower threshold of -60 dB is an artificial lower bound for calculating the hit list in our experiments. We present all of our results as *receiver-operating characteristics* (ROC) curves; hence, this lower bound can be thought of as the limit for computing the ROC. In practice, a much higher threshold would be chosen for an operational detector. The resulting hit locations are illustrated in figure 4.




Figure 4. Example hits from the pre-screener, where x's are FAs, +s are fiducial alarms and an * indicates a target.

In a final step, hits from the pre-screener are relocated to a local extrema for a user defined window in the raw FLGPR image. All of the hit locations are then sent to the feature extractor where the spectral analysis is computed on the original FLGPR image at those locations.

4. SPECTRAL FEATURES

To classify areas of interest in the radar imagery, features are extracted that can be used to model a positive detection class and a FA class. It is important to find features that produce separation between these two classes. As the results of this paper show, spectral characteristics of background clutter differs from that of explosive hazards. This difference is exploited to train a classifier and eventually classify new data. Figure 2 highlights the difference between the spectral responses of a target and a FA.

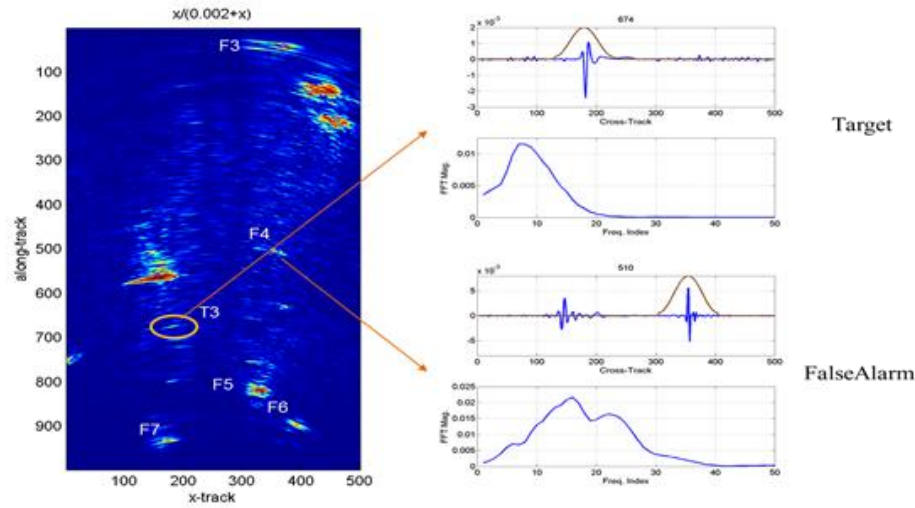


Figure 5. Illustration of spectral features in the cross track direction and the difference between actual targets and FAs.

The second and fourth rows in the second column of Figure 5 are the spectral responses of a target and a FA respectively. As this figure illustrates, there tends to be more randomness associated with a FA then with a target. There are three different forms of spectral features explored in this work, a 1D *fast Fourier transform* (FFT) in the cross track direction, multiple 1D cross track FFTs in the down track direction, and a 2-dimensional FFT.

4.1 One-Dimensional FFT Feature in the Cross Track Direction

The radar data contains complex images from which we use only the real portion in the spectral analysis. At each hit location provided by the pre-screener, a 1D FFT in the cross track direction is calculated. This is accomplished by taking a 101 point horizontal slice of pixels centered on the hit. A 101 point Hamming window is applied to this image slice so that more importance is given to the center of the alarm. The FFT is then performed on the row vector. The spectral features are obtained by computing the magnitudes of each frequency bin. A mathematical description of the process is

$$X_f(A) = |FFT(Re(H * I(A)\{-50:50\}))|,$$

where Re indicates that the real part of the complex image is used, H is the Hamming window, I is the windowed radar image and X is the resulting spectrum. We then store the 50 positive frequency elements of $X_f(A)$. These 50 frequency bins are the features used by the classifier. Figure 5 illustrates the output of the 1D FFT.

4.2 Multiple 1D FFT Cross Track Direction Features Taken in the Down Track Direction

The following approach requires calculating multiple 1D FFT slices. The FFT is calculated on multiple cross track slices of the radar image in the down track direction. Because the alarms are larger than one horizontal slice, using multiple slices will encompass more of the alarm and the surrounding area. This provides a more complete representation of the alarm compared to using a single slice of the image.

Eleven rows of 101 pixels, five above the alarm, five below and one centered on the alarm, are used in the calculation of the spectral features. Again a 101 point Hamming window is used on each of the slices and a 1D FFT is calculated for each of the eleven vectors. Figure 6 illustrates this concept.




Figure 6. Illustration of the multiple 1-D slice spectral features. Three cross track, from the eleven, 1-D features are extracted from a hit and shown.

Figure 6 shows three of these features, one centered on the target, one above center, and one below center. As with the 1D case, only the magnitude's of the first 50 frequency bins are used from each FFT, which are then concatenated to produce a feature vector with 550 elements.

4.2 Two-Dimensional FFT Feature

Although the multi-line FFT encompasses more of the target signature, it does not contain any down-track (vertical/columns) information. The 2D FFT contains both cross-track and down-track information, which provides a more complete spectral description of hits. Although the 2D and multi-line FFT use more of the image to extract features, our results show that these three feature sets can be used to complement each other.

To implement the 2D FFT, an image block, 101 pixels wide and 21 pixels high centered on the hit location, is used in the calculation. A Hamming window of the same size is applied to the image block to emphasize the center of the hit location. A 2D FFT is applied and the magnitude is found resulting in a spectral image that is 101 (columns) x 21 (rows). The equation is

$$X_f(A) = |FFT2(Re(H * I(A)\{-10:10, -50:50\}))|,$$

where H is a 2D Hamming window, I is a rectangular sub-image, and X is the 2D spectral image. The $FFT2$ notation indicates a 2D FFT operation. We only use the first quadrant of the 2D FFT producing a spectral image block that is 50 (columns) x 10 (rows). A rectangular block, instead of a square block, was used due to the tendencies of the target to be wider cross-track then down-track. Figure 7 shows the 2D FFT of a target with the DC shifted to the center of the image.

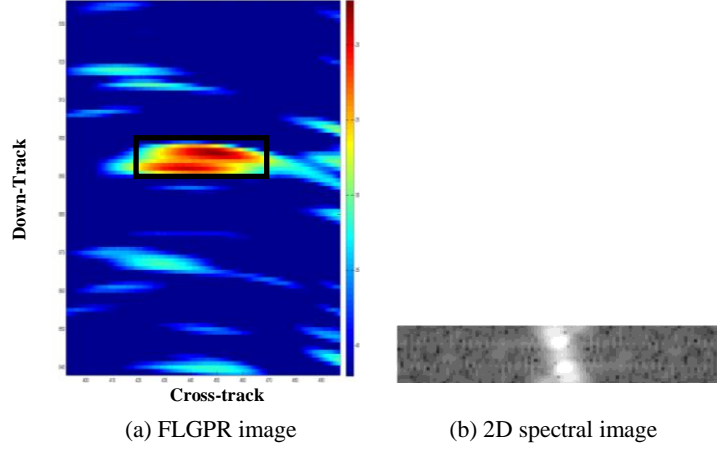


Figure 7. Illustration of the 2D FFT feature centered at a hit location.

5. CLASSIFIER


To reduce the number of false alarms, it is necessary discriminate between FAs and positive detections using the spectral features extracted from the radar imagery. To do this, we designed a classifier referred to as a genetic algorithm-based K-Nearest Neighbor probability density estimator (GA-based K-NN). This is a two-class classifier that uses the density of the training points surrounding a test point to calculate the probability of that test point belonging to each class. The training data we used had a greater number of FAs than positive detections, which means the density of FAs is much greater. To account for this, a K value is chosen for each class. A GA is used to learn these K values as well as reduce the size of the feature vector (known as feature selection).

5.1 GA-based K-NN

The training data are formed from the spectral features obtained from the radar imagery. When a new hit is indicated by the pre-screener, the spectral features are extracted from the radar for that area of interest. This new test point is then compared to the training elements for each class. The Euclidean distance to the K^{th} nearest neighbor is calculated, where K is different for each class. The K value for the FA class is typically much larger since there are many more FAs in the training set. This is exaggerated when the classes have a lot of overlap. On the other hand, if the data set has less overlap or is less lopsided, the K values may be very similar. These K values are learned by the GA and held constant once suitable values are determined. Doing several tests with the GA, we saw as much as 12% improvement in classification performance when different Ks were used, compared to having the Ks being the same for each class. Once the distance to the K^{th} training point is found, a probability is calculated using

$$P_C = \frac{K_C}{V * n_c},$$

where V is the volume of the hyper-sphere centered on the test point with a radius equal to the distance to the K^{th} training point and n is the number of training points for class C . Bayes rule is used to indicate the class to which the test point belongs. Figure 8 illustrates the process using a synthetic dataset. It demonstrates the added value of using different K values for each class.



Example 1:
 Two classes: Red, Blue
 100 Red points, 40 Blue points
 $K_{Red} = 9$
 $K_{Blue} = 3$
 $A_{Red} = \pi \times 0.22^2 = 0.16$
 $A_{Blue} = \pi \times 0.2^2 = 0.13$
 $P_{Red} = 3 / 100 / 0.16 = 0.56$
 $P_{Blue} = 3 / 40 / 0.13 = 0.58$
 $P_{Blue} > P_{Red}$, **Class Blue**

Example 2:
 Two classes: Red, Blue
 100 Red points, 40 Blue points
 $K_{Red} = 3$
 $K_{Blue} = 3$
 $A_{Red} = \pi \times 0.1^2 = 0.03$
 $A_{Blue} = \pi \times 0.2^2 = 0.13$
 $P_{Red} = 3 / 100 / 0.032 = 0.93$
 $P_{Blue} = 3 / 40 / 0.13 = 0.58$
 $P_{Blue} < P_{Red}$, **Class Red**

Figure 8. Two examples illustrating the calculations of the proposed classifier.

For the experiment shown in Figure 8, we generated two datasets from a normal distribution. The blue class has a mean of [0.3, 0.3], a covariance of [0.15, 0; 0, 0.15], and 40 points, while the red class has a mean of [0.75, 0.75], a covariance of [0.1, 0; 0, 0.1] and 100 points. This is typical of the data set from the FLGPR in that the positive detections tend to have a greater spread than the FAs and there are many more FAs. The test point is at [0.35, 0.35], which is very near to the mean of the blue class. Two images are shown as well as two results. The first image and result is obtained from having different values for K, while the second image and result is when the K values are the same. With this type of overlapping dataset, the result is typically better when the K values are different as evidenced by the examples in Figure 8. Since the test point actually belongs to the blue dataset, having the K values the same caused a misclassification. These K values can be chosen by the user, but it is much more practical to learn them with a GA.

5.2 Genetic Algorithm

The GA introduced in this work has two purposes. First, as mentioned above, it is necessary to learn the values of K for each class. Second, the GA is used to reduce the size of the feature vector. These objectives are done at the same time within the GA since each has an influence on the result of the other, meaning that the values for K will be different depending on what features are used.

We have found that it is not necessary to use all the features in order to model the data accurately. In fact, classification was improved when fewer features were used. In our previous work [2], we did an exhaustive search to find the best 4 out of 50 features from the 1D spectral analysis. This however was slow and time consuming when modifications were made to the other algorithm components, as each change in the components required retraining exhaustively. The GA provided a more efficient solution to this problem of feature selection. Another advantage is that the number of features necessary to build a good model can be learned as well. This is accomplished by using a variable size chromosome, where each gene in the chromosome is an integer (between 1 and 50 for the 1D spectral features) that represents a feature. A separate chromosome, with two genes, is used for the K values since they require a greater range of integers. We use a population of 50 chromosomes, each randomly generated to start. All the chromosomes are sent to the GA based K-NN, where the classifier uses them as parameters to classify the test data. Each chromosome pair (features and Ks), via the classifier, produces a ROC curve. The intention is to minimize the area above the ROC curve, therefore pushing the curve to the left. Figure 9 illustrates a ROC curve and the area we are trying to minimize highlighted in red.




Figure 9. Illustration of our goal to minimize the area above the ROC curve.

The GA uses the area above the ROC curve as a fitness for the chromosome pair. The smaller the area, the better the fitness value and the more fit the chromosome is. A new generation of chromosomes is made by roulette wheel selection, cross-over, and mutation operations. The roulette wheel selection chooses an intermediate population of 50 chromosome pairs. The chromosomes with greater fitness values are more likely to be chosen. These chromosomes are mated, based on the probability of cross-over (P_{co}), which we chose to be 85%. A single cross-over point is chosen at random and the mating is done, as in the example below.

Parent X	X1	X2	X3	X4	X5	X6	X7	K _{x1}	K _{x2}
Parent Y	Y1	Y2	Y3	Y4	Y5	K _{y1}	K _{y2}		
Child X	X1	X2	X3	Y4	Y5	K _{x1}	K _{x2}		
Child Y	Y1	Y2	Y3	X4	X5	X6	X7	K _{y1}	K _{y2}

Figure 10. Illustration of the single point cross-over with the chromosome pair.

Notice that the K values are not involved in the cross-over. The Ks that are associated with parent X also belong to child X and likewise with the Ks associated with parent Y. If a pair of chromosomes are not chosen for cross-over then the parents themselves become the children. Repeated genes in a chromosome (which may occur during cross-over) are removed, thus allowing the size of the chromosome to mutate. The possibility of a chromosome to be mutated is determined by the mutation rate, which we set to 10%. Mutation randomly chooses a gene and replaces it with a

randomly chosen number that isn't already in the chromosome, ensuring the uniqueness of the genes. Mutation is the only operation performed on the K values.

A copy of the best fit chromosome pair from the previous population automatically makes it into the next population with no modifications (often called "elitism"). The resulting children from cross-over and mutation become the new population and are then tested for their fitness. This cycle continues until an appropriate stopping criterion is met, specifically if the best fitness value hasn't improved over the course of seven generations.

Along with improved results, reducing the size of the feature vector provides faster testing time and more efficient operational performance. In our experiments, the number of 1D spectral features was typically reduced to 3 or 4, much less than the full set size of 50. For the multi-line spectral features we observed a reduction to around 30 features from 550. The 500 features produced by the 2-D FFT are often condensed to around 20 features.

6. RESULTS

6.1 FLGPR

Much improvement has been made in recent years on FLGPRs, yet it is still a maturing area of research. The ALARIC system used to collect the data for our research is still under development. Currently, more field tests are being conducted to improve the system. Our algorithms were developed using the limited amount of data so far collected by ALARIC. These data have given us a chance to refine our detection algorithms and learn the processing necessary to detect explosive hazards from radar. Specifically, our research into the use of spectral features for classification has proven successful even with this limited amount of data. Figure 11 shows preliminary results for the proposed classifier on confidence values for targets and FAs compared to the confidences from the pre-screener.

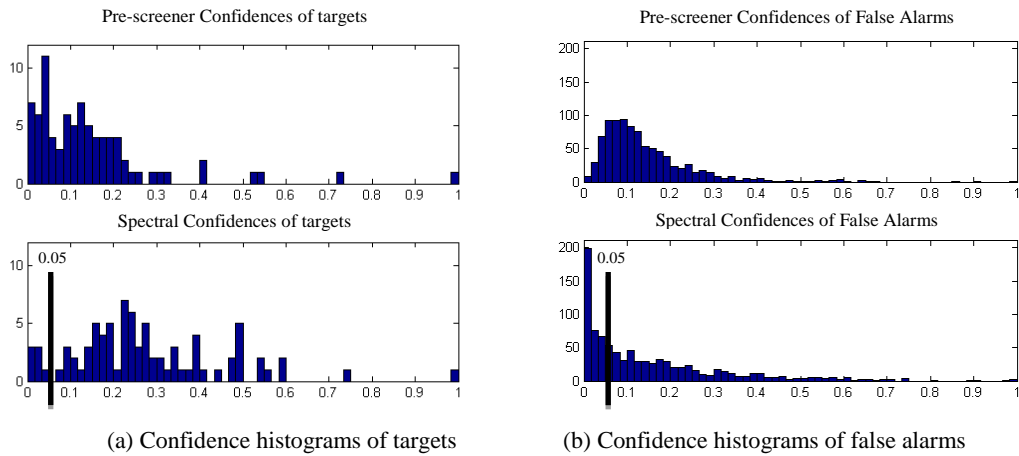


Figure 11. Confidence histograms for targets and false alarms from the pre-screener and the classifier.

For actual targets, the classifier confidences are shifted right, thus showing an increase in confidence values. At the same time, the confidences for FAs are shifted left, indicating a decrease in confidence values. This is desirable since we are working to make FAs have a low confidence and targets have a high confidence. Through visual inspection, if we set a threshold at 0.05 such that every confidence above 0.05 was called a target and everything below was called a FA, a significant number of FAs would be rejected.

These experiments were done using one lane for testing (lane1) and three lanes for training (lane 2, lane 3, and lane 4). There were 134 targets in the training set and 44 targets in the testing set. The GA was used to learn the features and the K values. Each chromosome was sent to the classifier which used lane 2, lane 3, and lane 4 as training to produce a ROC curve for lane 1.

This separation in confidences allows for better discrimination between FAs and targets. These preliminary results show that the spectral characteristics of FAs differ from that of explosive hazards. They also demonstrate that only a subset of features is necessary to model the data.

6.2 Confidence Maps

Results for FLGPR are preliminary at the moment due to the reality of working with a system that is currently under development. In this section we illustrate the proposed classifier and various spectral features for a different information source. Figure 12 shows the input source used in the experiment.




Figure 12. Confidence map using the MSER algorithm on LWIR.

In [4], we present an algorithm to detect buried explosive materials in LWIR using the maximally stable extremal regions (MSER) algorithm. The MSER is just one algorithm in our fusion system. The MSER is not currently being used in our fusion approach in a standalone fashion as a direct detection algorithm. The MSER does not assume any target shape, it just searches for spatially coherent connected regions. In our fusion article, we segment MSER regions in each LWIR image. Next, these image regions are projected into a quantized UTM map. We force each UTM map location to be observed in multiple images (a simple form of temporal aggregation of multiple LWIR images). For each MSER region in a LWIR image, the corresponding LWIR image values are projected to the UTM map and the final value at each location in the UTM map is the average of the multiple looks.

The following ROC curves are produced by varying the operating level (threshold) from the UTM map. At each operating level, we use the connected components algorithm to segment islands. The center of each island is a hit. A 1 meter halo is used. Any hit within 1 meter of a known target is lumped in with that ground truth location. Spectral features were collected for all the remaining hits and were then classified by the GA-based K-NN. The features and the K values for the classifier were learned as described above. These data have much more separation between the classes, allowing the K values to be very similar when learned by the GA. Since only one lane had the confidence map completed we used re-substitution (the testing and training data are the same) to obtain the results in Figure 13.

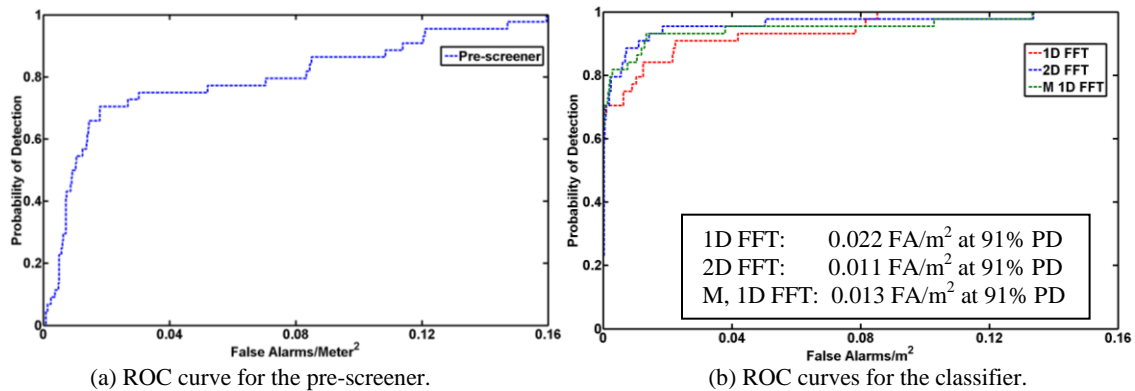


Figure 13. ROC curves for the pre-screener and the classifier.

In Figure 13(a) the ROC curves show the results of just the pre-screener. Figure 13(b) shows the results of the classifier using the 1D, 2D and multi-line FFT spectral features. A significant improvement in performance is observed. Although

these ROC curves reach 100% PD, the MSER algorithm only detected 44 out of the 50 targets in the lane. The ROC curves represent a detection of 100% of the 44 targets the MSER found.

The results show the substantial improvement obtainable by using spectral features. The pre-screener was only capable of 80% PD at about 0.08 FA/m². A significant improvement was made using the spectral features which had a PD of 91% at 0.011 FA/m² for the 2D FFT. Furthermore, the 2D FFT provided much better results than the 1D FFT, which isn't surprising because the 2D FFT encompasses more of the radar image around the hit location. The Multi-line FFT also showed better performance than the 1D FFT and performs almost as well as the 2D FFT. Through our experiments we observed that the 1D FFT, the multi-line FFT, and the 2D FFT tend to detect different targets at earlier points on the ROC curve. By combining the detections of the classifier for all three feature sets at 90% PD we were able to detect 43 out of the 44 targets (97.7% PD). This shows that these feature sets complement each other and when used together they can provide improved PD and FA rates.

7. Summary

The framework we proposed has shown to provide an effective means of detecting explosive hazards. The LASDF pre-screener is a significant improvement over a non-adaptive approach, producing an improved hit list without an unreasonable amount of FAs. Spectral features have shown to be successful in discriminating between background clutter and positive detections. The extension into multi-line FFT and the 2D FFT features provided a more comprehensive set of features with which to classify hits. Our empirical results show that the 2D FFT was able to decrease the FA rate by 50% at >90% probability of detection compared to the 1D FFT.

The GA-based K-NN, along with the spectral features, allowed us to build a model of the FAs and the targets for a two-class classifier. On average, a greater than 12% improvement in the area above the ROC curves was achieved using different K values for each class with the radar data. On the other hand with data that shows less overlap between classes, as shown in the confidence maps, the K values tended to be more similar. Having this flexibility for the K values allows our algorithm to accommodate various types of datasets.

Finally, the GA was useful in determining the K values for the classifier as well as reducing the size of the feature vector significantly. The smaller feature vector decreased the amount of computation time by 4x on average when compared to using all the features. This reduced feature set also allowed for a more generalized model of the data, producing better results. Overall, this framework has shown to be effective and, as advancements are made with FLGPR, we believe the tools proposed here will be valuable for the detection of explosive hazards.

ACKNOWLEDGEMENTS

This work was funded by Army Research Office grant number 57940-EV to support the U. S. Army RDECOM CERDEC NVESD, and by Leonard Wood Institute grant LWI 101-022. The authors would like to thank Dr. Don Reago, Mr. Pete Howard, and Mr. Richard Weaver at NVESD and Dr. Russell Harmon at ARO.

REFERENCES

- [1] Ho, K.C., Gader, P.D., "Improving landmine detection using frequency domain features from ground penetrating radar", Proc. SPIE 5794, 1141-1150 (2005).
- [2] Havens, T.C., Ho, Farrell, J., K.C., Keller, J.M., Ton, T.T., Wong, D.C., and Soumekh, M., "Locally-adaptive detection algorithm for forward-looking ground penetrating radar", Proc. SPIE 7664, (2010).
- [3] Havens, T.C., Spain, C.J., Ho, K.C., Keller, J.M., Ton, T.T., Wong, D.C., and Soumekh, M., "Improved Detection and False Alarm Rejection Using FLGPR and Color Imagery in a Forward-Looking System", Proc. SPIE 7664, (2010).
- [4] Anderson, D.T., Keller, J.M., Sjahputera, O., "Algorithm fusion in forward looking long wave infrared imagery for buried explosive hazard detection", Proc. SPIE, (2011).